



Computer Networks - Midterm

November 3, 2017

Duration: 2h15m

- This is a closed-book exam.
- Please write your answers on these sheets in a readable way, in English or in French.
- You can use extra sheets if necessary (don't forget to put your name on them).
- The total number of points is 100.
- This document contains 21 pages.
- Good luck!

Full Name (Nom et Prénom):

SCIPER No:

Division: Communication Systems Computer Science
 Other (mention it):

Year: Bachelor Year 2 Bachelor Year 3
 Other (mention it):

1 Short questions

(10 points)

For each question, please circle a single best answer.

1. Caching of DNS responses can:
 - (a) speed up hostname-to-IP address resolution.
 - (b) save bandwidth.
 - (c) improve resistance to flooding attacks.
 - (d) All of the above.
2. Compared to a packet switched network, a connection switched network:
 - (a) has higher packet loss and delay.
 - (b) does better use of bandwidth.
 - (c) uses resources less efficiently.
 - (d) None of the above.
3. Which of the following is true about web cookies?
 - (a) A cookie is state created by the web client.
 - (b) Cookies speed up the delivery of web content.
 - (c) Cookies allow web sites to tailor what the user sees on the screen.
 - (d) None of the above.
4. How many iterative DNS queries a local DNS server sends when trying to resolve `nal.epfl.ch`?
 - (a) At most 1.
 - (b) Exactly 3.
 - (c) At most 3.
 - (d) More than 3.
5. Which of the following is true?
 - (a) A socket enables a process to communicate with another process.
 - (b) A process can use a UDP socket to communicate with many remote processes.
 - (c) A process can use a TCP socket to communicate with only one remote process.
 - (d) All of the above.

6. A Distributed Hash Table (DHT) stores:
- (a) content files.
 - (b) only metadata files.
 - (c) only popular content.
 - (d) the identities of the peers that store each content file.
7. Suppose that we double the rate of all links in a network. Which of the following is true?
- (a) The end-to-end delay will decrease by at most 50%.
 - (b) The end-to-end delay will decrease exactly by 50%.
 - (c) The propagation delay will decrease exactly by 50%.
 - (d) The queuing delay will decrease exactly by 50%.
8. A user generates data as follows: 1% of the time he is active and produces data at 100 Mbps, while he is silent for the rest of the time. How many such users can a link of 1 Gbps support while guaranteeing no packet loss?
- (a) Less than 10.
 - (b) Exactly 10.
 - (c) More than 10.
 - (d) We cannot say precisely, all the above can be true.
9. For a 100 Kbps link with a 40 msec end-to-end propagation delay and packets of 8 Kbits, what is the sender utilization (the fraction of time the sender is busy), if the Stop-and-wait protocol is used and the size of the ACKs is 1 bit?
- (a) Less than 0.5.
 - (b) Exactly 0.5.
 - (c) More than 0.5.
 - (d) We cannot say precisely, all the above can be true.
10. Suppose a switch with an initially empty buffer of infinite capacity and an outgoing link with transmission rate 1Mbps. Each packet is of 100Kbits and batches of 10 of them simultaneously arrive at the switch every 1 sec. The average queuing delay a packet experiences is therefore:
- (a) 0.1 sec.
 - (b) between 0.1 and 1 sec.
 - (c) 1 sec.
 - (d) None of the above.

2 Web Browsing and DNS

(30 points)

Setup:

Three users, Alice, Bob, and Persa, are logged into their workstations, respectively called `alice.ethz.ch`, `bob.ethz.ch`, and `persa.ethz.ch`, all located inside ETHZ's network.

ETHZ offers a web server `www.ethz.ch` and a local name server `ns.ethz.ch`, which is also the authoritative server for the `ethz.ch` domain.

EPFL offers a web server `www.epfl.ch` and a local name server `ns.epfl.ch`, which is also the authoritative server for the `epfl.ch` domain.

The setup is illustrated in Figure 1. Make the following assumptions:

- All DNS caches are initially empty.
- All DNS requests are resolved recursively.
- Each HTTP message fits in a single packet.
- The web browsers and servers use persistent HTTP connections.

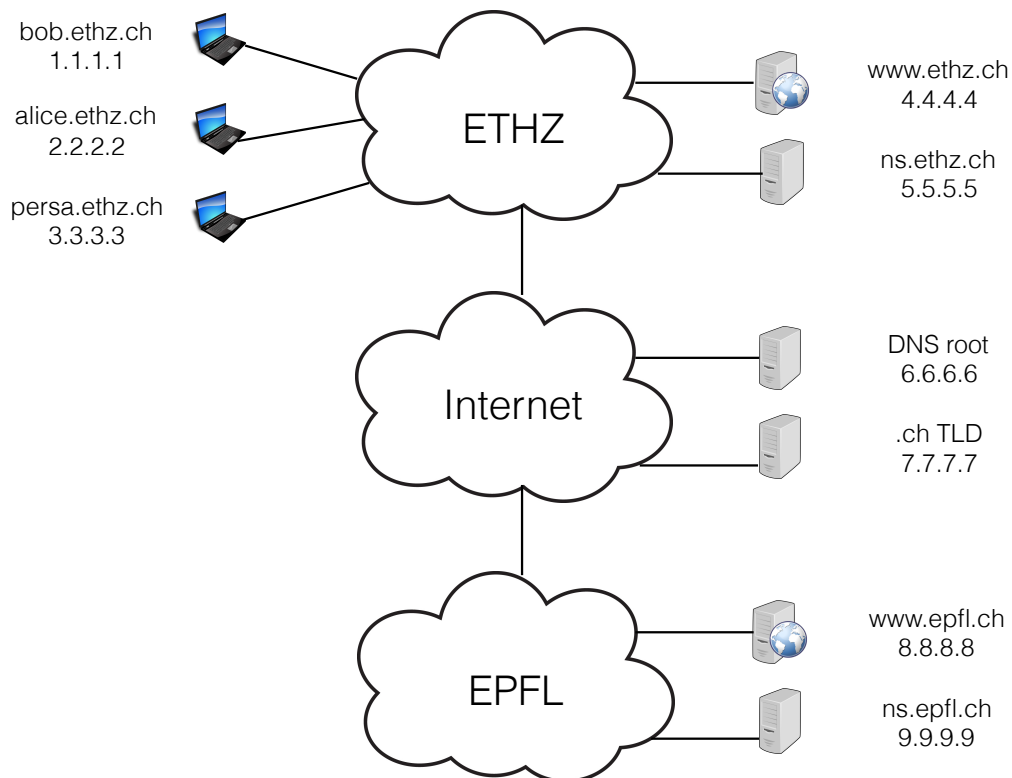


Figure 1: The setup for the exercise on Web browsing and DNS

Question 1 (10 points):

Alice uses her web browser to access URL `http://www.epfl.ch/index.html`. This page references one image with URL `http://www.epfl.ch/image.jpg` and another web page with URL `http://www.ethz.ch/file.html` (which does not reference any other object).

List all packets that are exchanged in the entire network, including any connection-setup packets, by filling Table 1. For each packet, show the source and destination IP address, the transport-layer protocol, the application-layer protocol, and the purpose of the packet, as in the example.

Packet	Source IP	Dest. IP	Transport protocol	Application protocol	Purpose
ex.1	1.0.0.1	1.0.0.2		HTTP	HTTP reply with image.jpg
1					
2					

Table 1: Packets transmitted on the Internet, in Question 1

Question 2 (5 points):

Which is the minimum number of sockets on Alice's machine that are necessary for the scenario in Question 1? What kind of sockets are they? Which of the packets that you listed in Table 1 are sent and received over each socket?

Question 3 (5 points):

After Alice has viewed `http://www.epfl.ch/index.html`, Bob wants to access the same URL.

Persa is a malicious user who guesses exactly when Bob tries to access `http://www.epfl.ch/index.html`. She wants to trick Bob and make him access a web server running on her own workstation, thinking that he is accessing the EPFL web server.

How can Persa do that by sending DNS traffic to Bob?

Question 4 (10 points):

Describe in more detail Persa's attack from Question 3:

a) Describe the packet(s) sent by Persa. What will be the source IP address and port number, destination IP address and port number, transport-layer protocol, and application-layer protocol of each packet?

This is a relatively hard question. Unless you have already thought about this attack and know the answer, we suggest leaving it for the end.

3 Reliable Data Transport

(30 points)

Question 1 (6 points):

End-host A wants to send a file to end-host B using a transport-layer protocol that provides reliable data delivery with pipelining and sender window $N=5$. A's transport layer splits the file in 10 segments with sequence numbers from 0 to 9. **The first data segment from A to B is lost the first time it is transmitted**, while all other transmissions arrive at their destination uncorrupted and in order.

Complete the two diagrams in Figure 2 to show all the segments (with sequence numbers and ACK numbers) transmitted by A and B, in each of the following two cases:

- A and B use Go-Back-N (left part of Figure 2).
- A and B use Selective Repeat (right part of Figure 2).




Figure 2: Sequence diagram to be completed for Question 1.

Question 2 (6 points):

Consider the same scenario as in Question 1, with the following difference: **The first 4 ACK segments from B to A are lost**, while all other segments arrive at their destination uncorrupted and in order.

Complete the two diagrams in Figure 3 to show all the segments (with sequence numbers and ACK numbers) transmitted by A and B, in each of the following two cases:

- A and B use Go-Back-N (left part of Figure 3).
- A and B use Selective Repeat (right part of Figure 3).




Figure 3: Sequence diagram to be completed for Question 2.

Question 3 (3 points):

Based on your answers to Questions 1 and 2, can you draw a general conclusion about when one should prefer Go-Back-N over Selective Repeat, and when one should prefer Selective Repeat over Go-Back-N? Justify your answer, e.g., if you state that one should prefer Go-Back-N in situation X, please say what metric Go-Back-N improves relative to Selective Repeat in situation X.

Question 4 (15 points):

In class, we discussed different elements of reliable data transfer (RDT): checksums, acknowledgments, timeouts, sequence numbers, and pipelining.

End-host A wants to send data reliably to end-host B (B never sends data to A). Which of the above RDT elements do they need to use in each of the following scenarios and why? In all scenarios except for (e), A wants to send data to B with high throughput.

- a) The path from A to B may corrupt packets, but never drops or reorders packets. The path from B to A never corrupts, drops or reorders packets. There are no queuing/processing delays on either path.

b) Both paths (from A to B, and from B to A) may corrupt packets, but they never drop or reorder packets, and they never introduce queuing/processing delays.

c) Both paths (from A to B, and from B to A) may drop or reorder packets or introduce queuing/processing delays, but they never corrupt packets.

d) The path from A to B never corrupts, drops or reorders packets, and never introduces queuing/processing delays. The path from B to A may do any of these things.

e) A sends one small piece of data (small enough to fit in one segment) to B every hour. Both paths (from A to B, and from B to A) may corrupt, drop or reorder packets or introduce queuing/processing delays.

4 Delay Computation

(30 points)

The network in Figure 4 contains the following elements:

- Server S , which stores a file.
- N “odd” nodes and N “even” nodes, where $N \geq 2$.
- Packet switches $S1$, $S2$ and $S3$, which perform store-and-forward packet switching and introduce insignificant processing delays. Each switch has a separate (infinite) queue for each outgoing link. E.g., switch $S2$ has three separate queues: one queue for packets addressed to the “odd” nodes, one queue for packets addressed to the “even” nodes, and one queue for packets addressed to server S . As soon as $S2$ receives a packet, it immediately places the packet in one of the three queues, depending on the packet’s destination.

Assume that:

- Each link has length ℓ and propagation speed c .
- The transmission rate of a link is the same in both directions of the link.
- The end-systems communicate over UDP. When you compute delays, you do not need to worry about connection setup.
- The end-systems communicate using the client-server architecture. There is no P2P file distribution in this problem.
- No other approximation/simplifying assumption can be made.

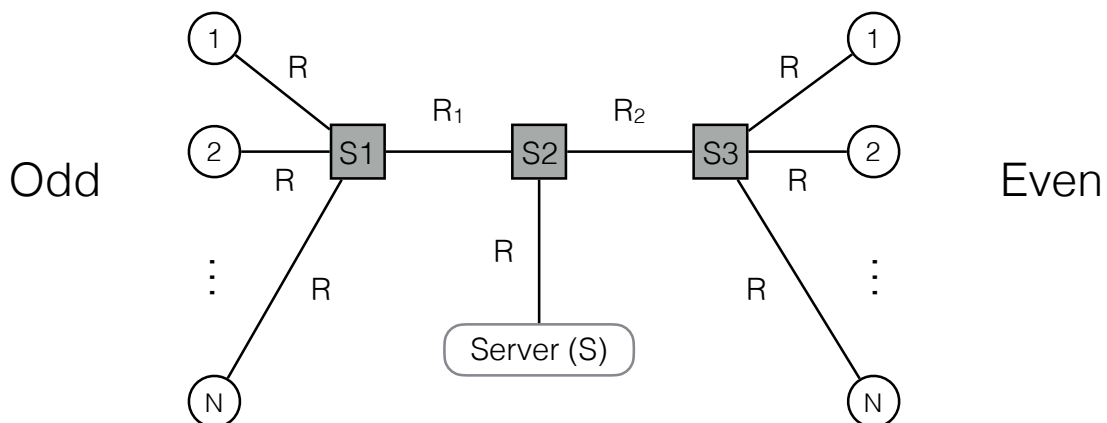


Figure 4: Network topology

Question 1 (15 points): Suppose that: $R_1 < R_2 < R$.

- a. Each “odd” node sends a request to server S for the file (no “even” node sends a request). All the “odd” nodes start transmitting their requests simultaneously. Each request fits in a single packet of size Q bits. How long does it take for all the “odd” requests to reach the server S ?

- b. Each “odd” and each “even” node has sent a request to server S for the file, and S has received all $2N$ requests. S splits the file into P packets of equal length L and sends a copy of the file to each node. S sends the file first to all the “even” and then to all the “odd” nodes.

How long does it take for the file to be fully delivered to all the “odd” and “even” nodes?

Question 2 (15 points): Now, suppose that: $R < R_1 < R_2$ and $NR_1 > R_2$.

- a. Each “odd” and each “even” node sends a request for the file to server S . All nodes start transmitting their requests simultaneously. Each request fits in a single packet of size Q bits.
- (i) Compute the time $T_{first.odd}$ at which the first “odd” request arrives at S_2 .
 - (ii) Compute the time $T_{last.even}$ at which the last “even” request arrives at S_2 .
 - (iii) Will the first “odd” request experience queuing delay at switch S_2 ?
 - (iv) How long does it take in total for all the “odd” and “even” requests to reach S ?

- b. Server S has received all $2N$ requests sent in Question 2b. S splits the file into P packets of equal size L and sends a copy to each node. S sends the file first to all the “even” nodes and then to all the “odd” nodes.

How long does it take for the file to be fully delivered to all the “odd” and “even” nodes?